HPSS for Spectrum Scale

Point-in-time backup/restore and automatic space management for Spectrum Scale

Spectrum Scale for HPSS overview

- Not only is HPSS a highly scalable standalone file repository, a single HPSS may also provide disaster recovery and space management services for one or more Spectrum Scale file systems.
- Spectrum Scale customers may now store petabytes of data on a file system with terabytes of high performance disks.
- HPSS may also be used to backup your Spectrum Scale file system, and when a catastrophic failure occurs, HPSS may be used to restore your cluster and file systems.

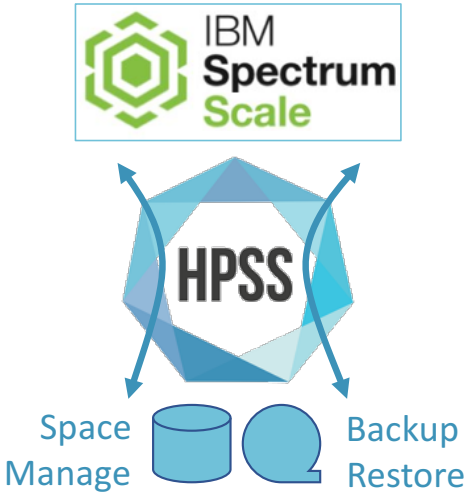

Many-to-one advantage of HPSS for Spectrum Scale

- A single HPSS may be used to manage one or more Spectrum Scale clusters.
- ALL of the HPSS storage may be shared by all Spectrum Scale file systems.
- One Spectrum Scale file system may leverage all HPSS storage if required.
- The Max Planck Computing and Data Facility (MPCDF, formerly known as RZG) is space managing and capturing point-in-time backups for six (6) Spectrum Scale file systems with ONE HPSS.

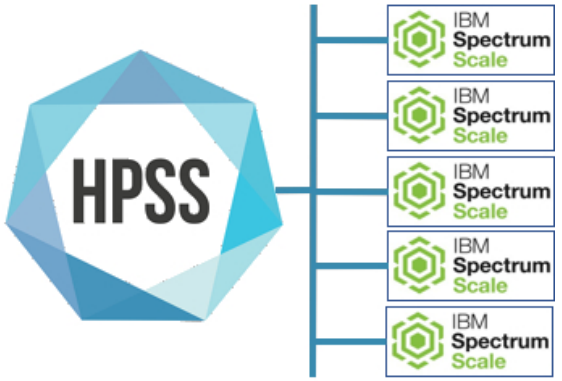

HPSS software for Spectrum Scale

- HPSS Session software directs the space management and disaster recovery services.
- HPSS I/O Manager (IOM) software manages data movement between HPSS and Spectrum Scale.
- HPSS Session software is configured on all Spectrum Scale Quorum nodes.
  - Session software is only active on the CCM (Spectrum Scale Cluster Configuration Manager) node.
- The CCM may failover to any Quorum node
  - Session software will follow the CCM.
- HPSS IOM software may be configured to run on any node with a Spectrum Scale mount point.

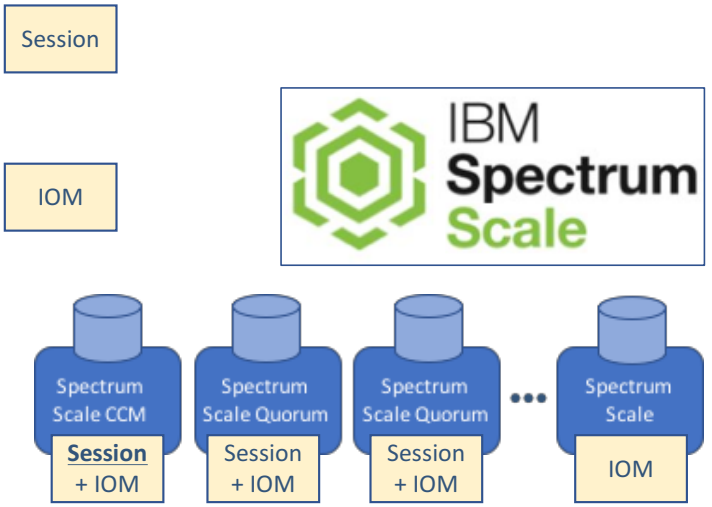

HPSS software for Spectrum Scale

- There are five processes that comprise the HPSS Session software:
  - HPSS ProcessManager
  - HPSS MountDaemon
  - HPSS ConfigurationManager
  - HPSS ScheduleDaemon
  - HPSS EventDaemon
- There are three processes that comprise the distributed HPSS IO Manager software:
  - HPSS IOManager – manage I/O between HPSS and Spectrum Scale
  - HPSS IOAgent – high speed transfers of individual large files
  - HPSS ishTAR – high speed transfer of many small files

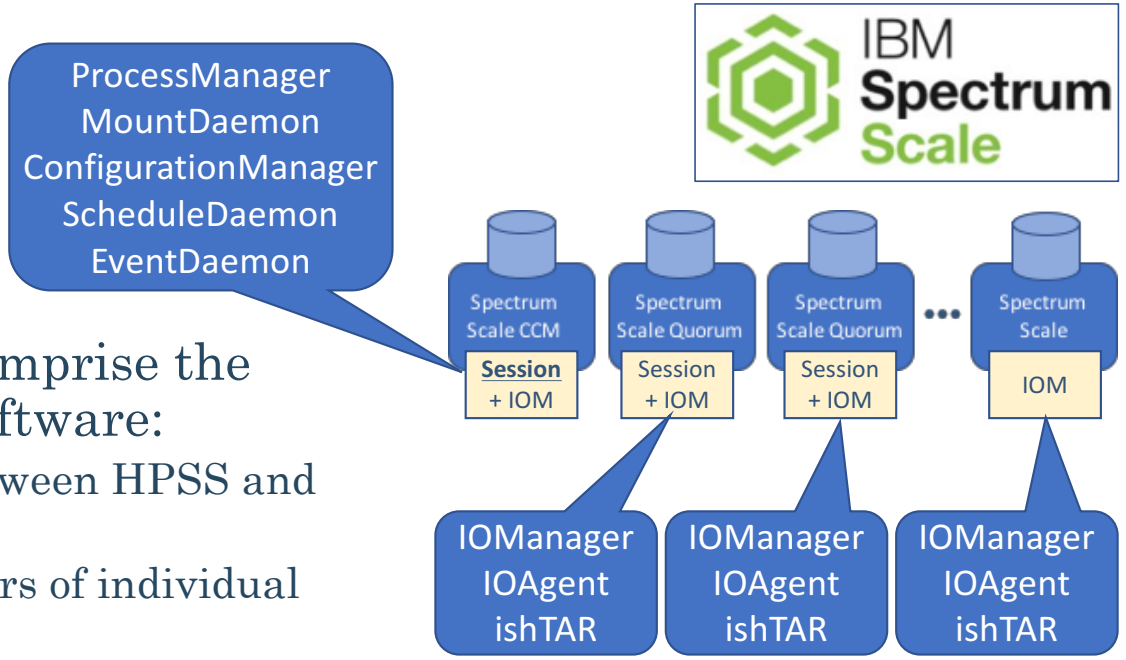

HPSS hardware for Spectrum Scale

- HPSS IOMs are highly scalable and multiple IOMs may be configured on multiple Spectrum Scale nodes for each file system.
- Bandwidth requirements help determine the expected server count for a deployment.
- New Spectrum Scale quorum nodes for HPSS Session and HPSS IOM software may need to be added to the cluster.

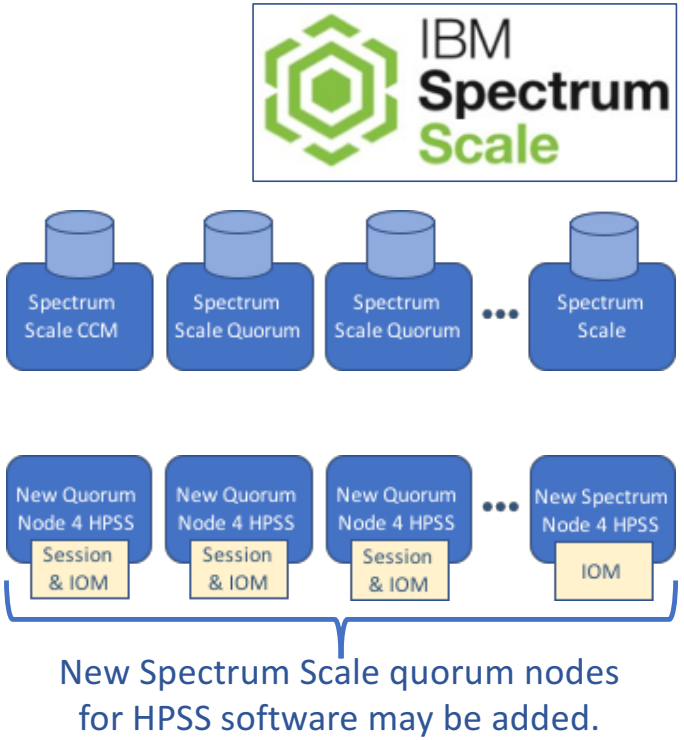

Space manage Spectrum Scale with HPSS

- Periodic ILM policy scans are initiated by the HPSS Session software to identify groups of files that must be copied to HPSS.
- The HPSS Session software distributes the work to the I/O Managers.
- Spectrum Scale data are copied to HPSS in parallel.
- Periodic HPSS copies files from Spectrum Scale to HPSS.

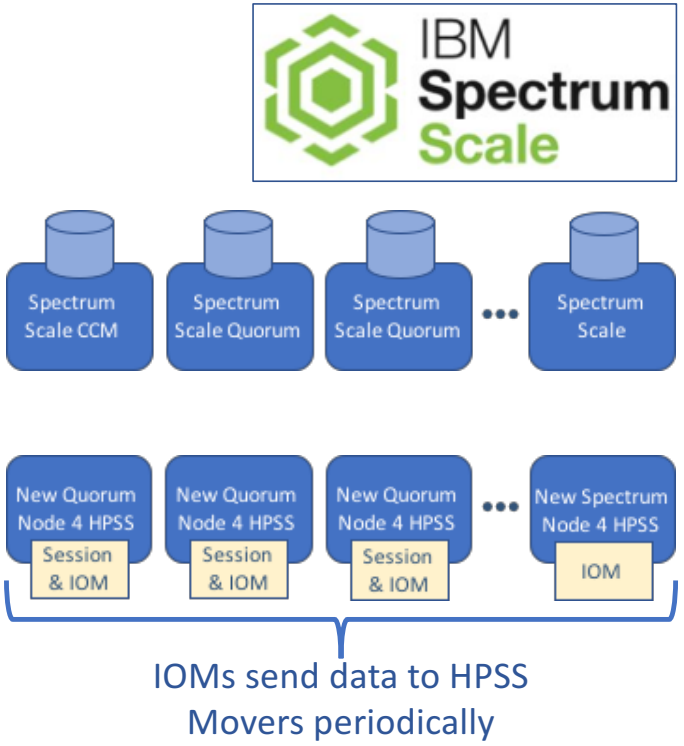

Space manage Spectrum Scale with HPSS

- When Spectrum Scale capacity thresholds are reached (the file system is running out of storage capacity), unused files are purged from Spectrum Scale, but the inode and other attributes are left behind.
- All file names are visible, and the user may easily identify which files are online and which files are near-line.
- The HPSS Session software will automatically recall any near-line file from HPSS back to Spectrum Scale when accessed.
- The ghi_stage command-line tool may be used to efficiently recalling a large number of files.

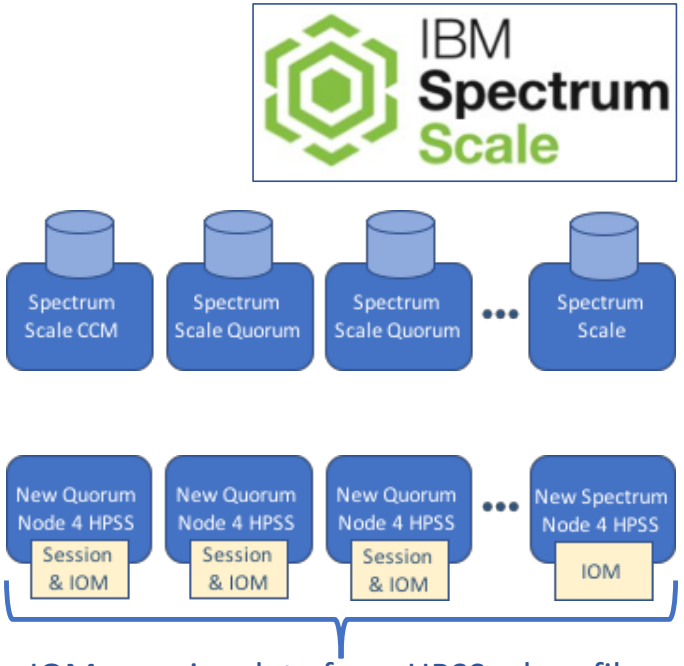

Backups & disaster recovery with HPSS

- The backup process captures the following data:
  - File data - the space management process (discussed above) is the process used by HPSS to capture the data for each file.
  - Namespace data - the Spectrum Scale name space is captured using the Spectrum Scale *image backup* command.
  - Cluster configuration data - the cluster configuration is saved to HPSS to protect the cluster configuration.
- The HPSS disaster recovery processing minimizes data movement and is ideal for high performance computing (HPC) environments where the goal is to bring the namespace back online quickly.